Of Data Dropboxes and Data Gloveboxes

Bloomberg

Engineering

Clay Baenziger
Product Owner - Hadoop Infrastructure

September 28, 2018

TechAtBloomberg.com

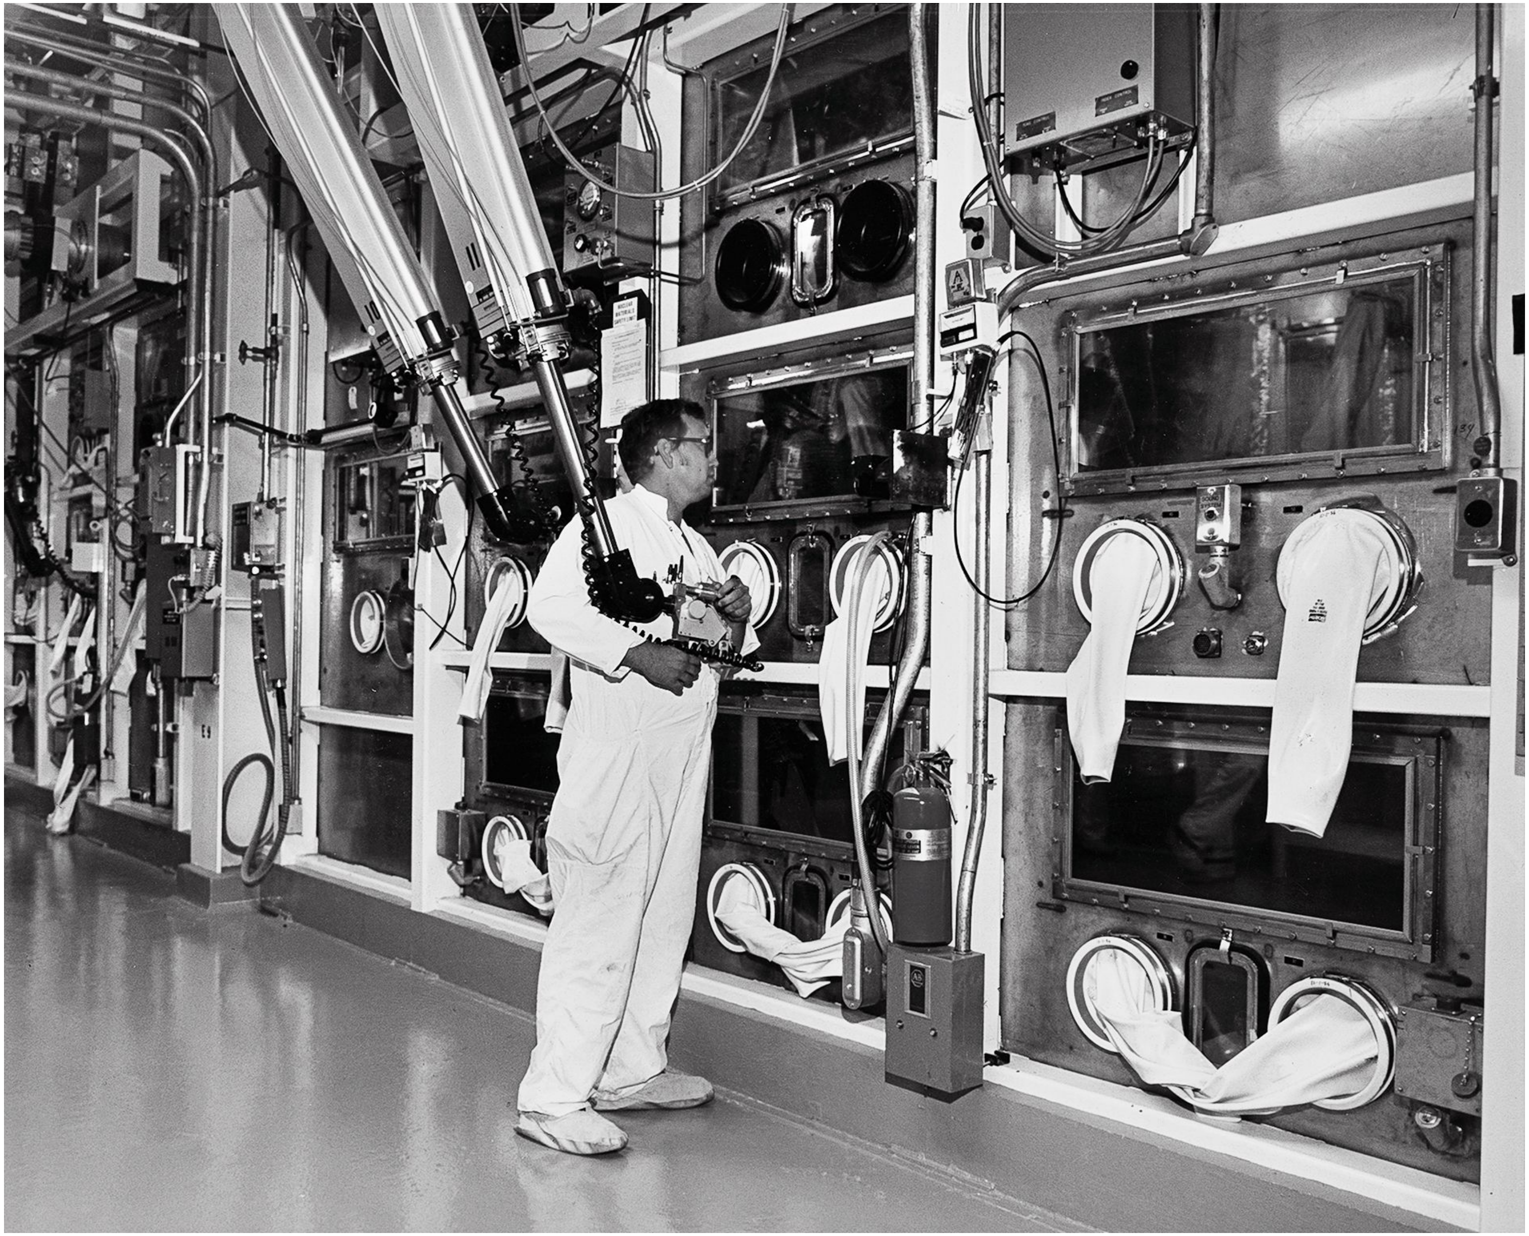

Image: Denver Public Library, Rocky Mountain News Photographic Archives, "Rocky Flats employee handles a robotic arm assembly." 11 Nov. 1987

Bloomberg By the Numbers

TechAtBloomberg.com

Bloomberg
Engineering

Bloomberg By the Numbers

- Founded in **1981**
- **325,000** subscribers in **170 countries**
- Over **19,000 employees** in 192 locations
  - **Over 5,000 software engineers**
  - 100+ machine learning data scientists and engineers
- **More News reporters** than The New York Times + Washington Post + Chicago Tribune
  - News content from 125K+ sources
  - >1.5M news stories ingested / published each day (that's 500 news stories ingested/second)
- One of the largest private networks in the world
- 100B+ tick messages per day, with a peak of more than 10 million messages/second
- More than a billion messages (E-Mails and IB chats) processed each day

Handling Material

Image: Office of Legacy Management, U.S. D.O.E., “[Rocky Flats Overview](#)”, 20 Aug. 2014. Pg. 13

Properties of Radioactive Materials

Radioactive Material:

- Can be harmful to people in small quantities
- Can have a very long hazard life if released
- Should be isolated to prevent their spread
- Should be cataloged and characterized to assess harm
- Can still be machined and worked with proper technique

— Robotics

— Personal Protective Equipment

Properties of Data

Data:

- Can be harmful to people in small quantities
- Can have a very long hazard life if released
- Should be isolated to prevent their spread
- Should be cataloged and characterized to assess harm
- Can still be used and analyzed with proper technique
  - Continuous/Automated Deployment
  - Workflow Automation and Gloveboxes

Nuclear Material (non)Proliferation

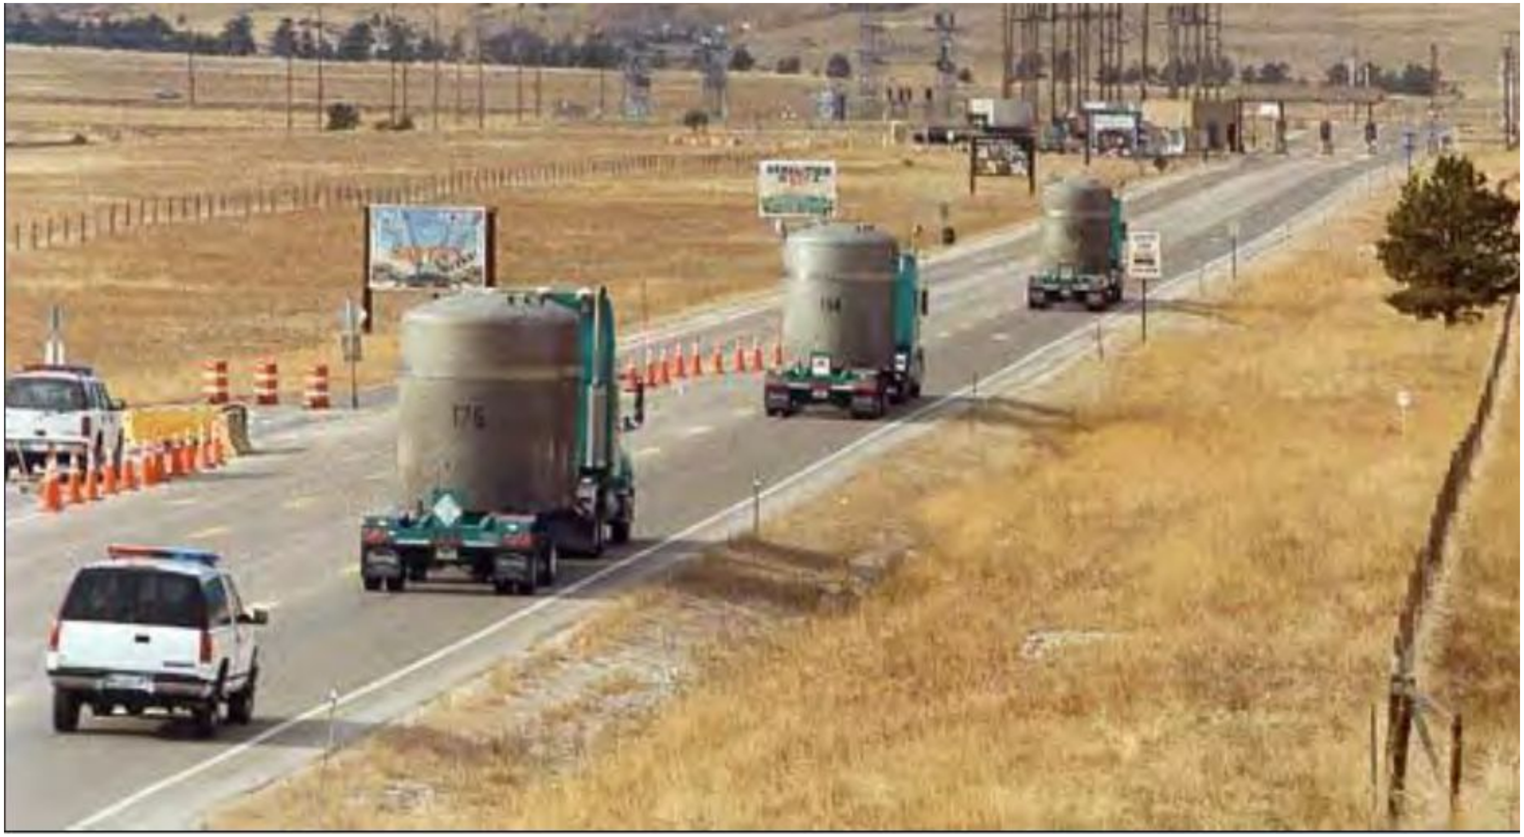

Image: Office of Legacy Management, U.S. D.O.E.,[”Rocky Flats Overview”](), 20 Aug. 2014. Pg. 59

Data Material Proliferation

- Terrorists? (Well certainly hackers...)
- Accidental loss (USB sticks, laptops, etc.)
- No Price-Anderson Act for Data Incidents
  - Quite the opposite with GDPR!
  - GDPR limits untraceable mixing of data
- Data Sovereignty
  - Requires data to remain geographically stationary
  - Must move computation to the data
- Data:
  - Swamps
  - Lineage
  - Masking

Clean-Up Is Messy

Image: CO Dept. of Pub. Health, "Citizen Summary Rocky Flats Historical Public Exposures Studies 1969 Fire",

Hadoop Cluster

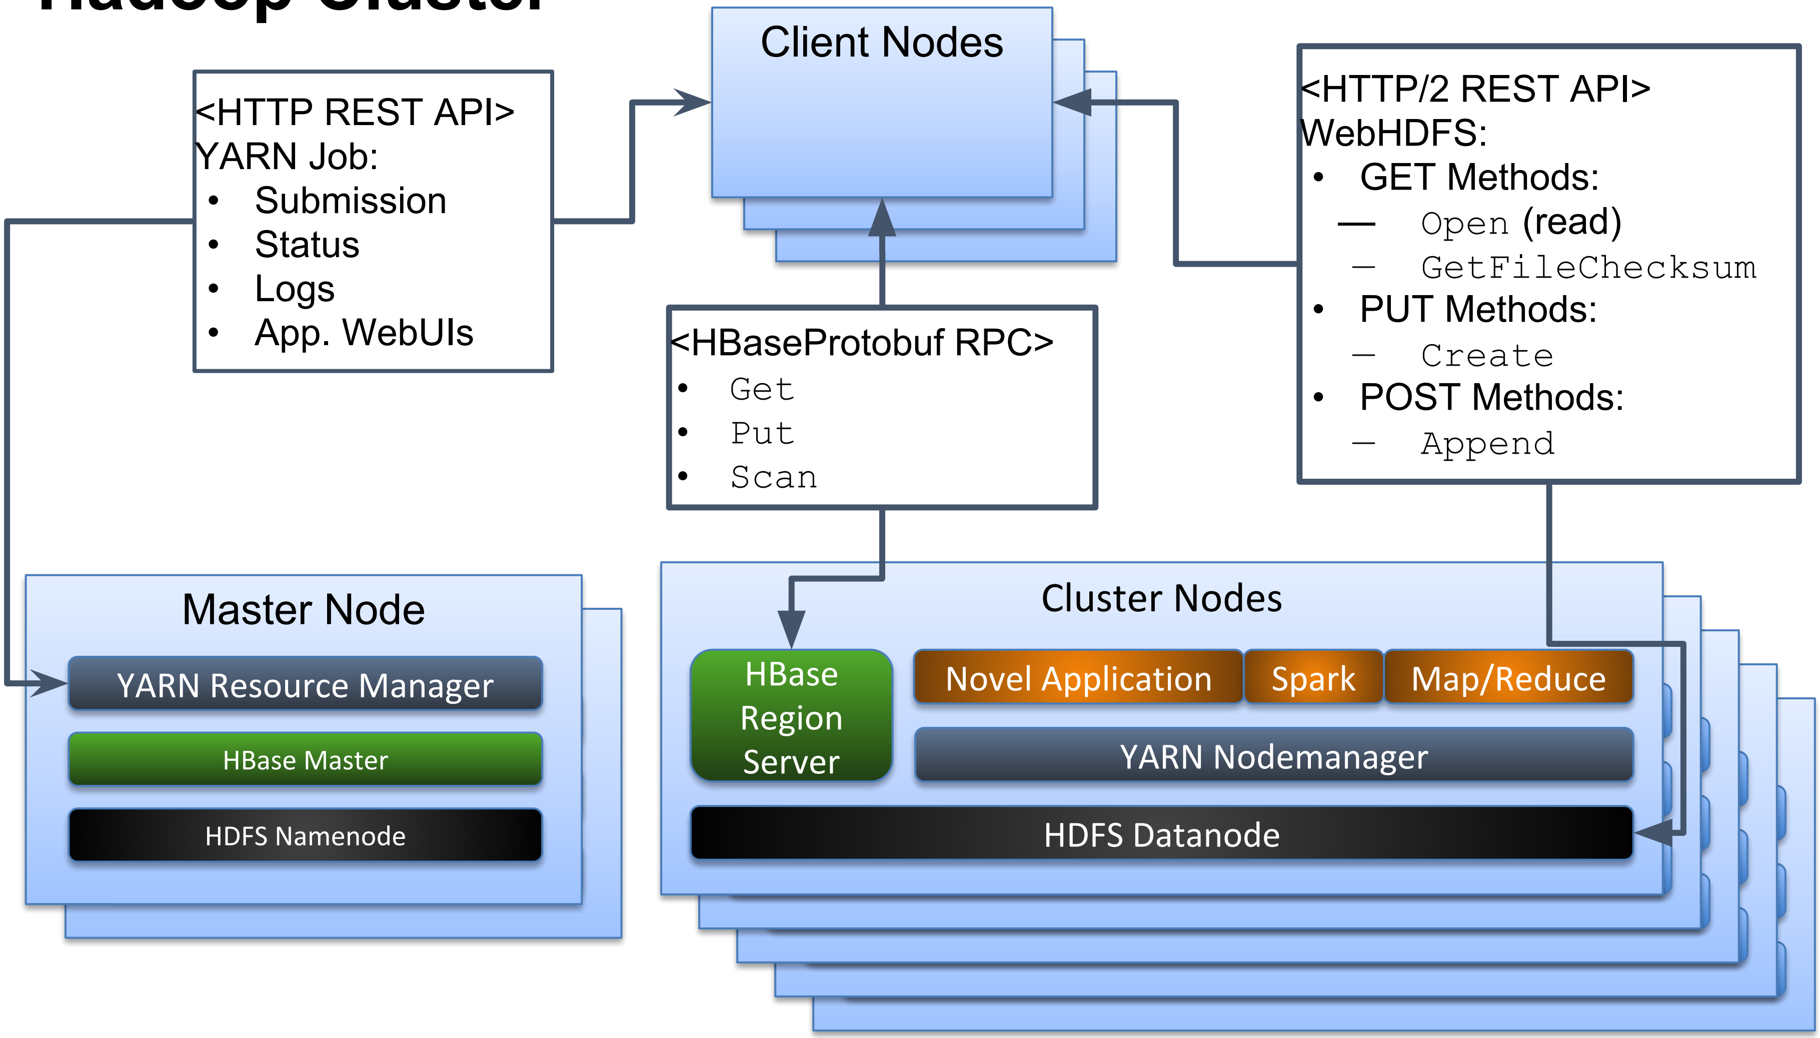

Client Nodes

<HTTP REST API>
YARN Job:
- Submission
- Status
- Logs
- App. WebUIs

<HBaseProtobuf RPC>
- `Get`
- `Put`
- `Scan`

<HTTP/2 REST API>
WebHDFS:
- GET Methods:
  - `Open` (read)
  - `GetFileChecksum`
- PUT Methods:
  - `Create`
- POST Methods:
  - `Append`

Master Node

YARN Resource Manager

HBase Master

HDFS Namenode

Cluster Nodes

HBase Region Server

Novel Application

Spark

Map/Reduce

YARN Nodemanager

HDFS Datanode

TechAtBloomberg.com

Bloomberg
Engineering

Lock Everything Down

Image: Office of Legacy Management, U.S. D.O.E.,["Rocky Flats Overview"](), 20 Aug. 2014. Pg. 21

Lock Everything Down

- Firewalls
- Encryption:
  - At rest
  - On the wire
- Authentication
- Limit Client Connections
- Dropboxes:
  - Data goes in
  - Data can not come out
  - Provide validation of transmission
  - At petabyte scale?

DropboxFilter for HDFS (WebHDFS API)

- Upload:

```
$curl -X PUT "http://<HOST>:<PORT>/webhdfs/v1/<PATH>?op=CREATE
HTTP/1.1 307 TEMPORARY_REDIRECT
$curl -X PUT -T <File> "http://<DN>:<PORT>/webhdfs/v1/<PATH>?op=CREATE..."
HTTP/1.1 201 Created
```

- Download:

```
$curl -L "http://<HOST>:<PORT>/webhdfs/v1/<PATH>?op=OPEN
```

- Checksum:

```
$curl -L "http://<HOST>:<PORT>/webhdfs/v1/<PATH>?op=GETFILECHECKSUM"
{"FileChecksum": {
  "algorithm": "MD5-of-0MD5-of-512CRC32C",
  "bytes":     "[...]00eb745ad2f5bd1dccab359b12f7f9411b00000000",
  "length":    28
}}
```

DropboxFilter for HDFS (Architecture)

- HDFS Protocols
  - Protobuf RPC
  - RESTful API over HTTPS/2
- Servlet Based Web Server Design
  - Filter
  - Request Handler

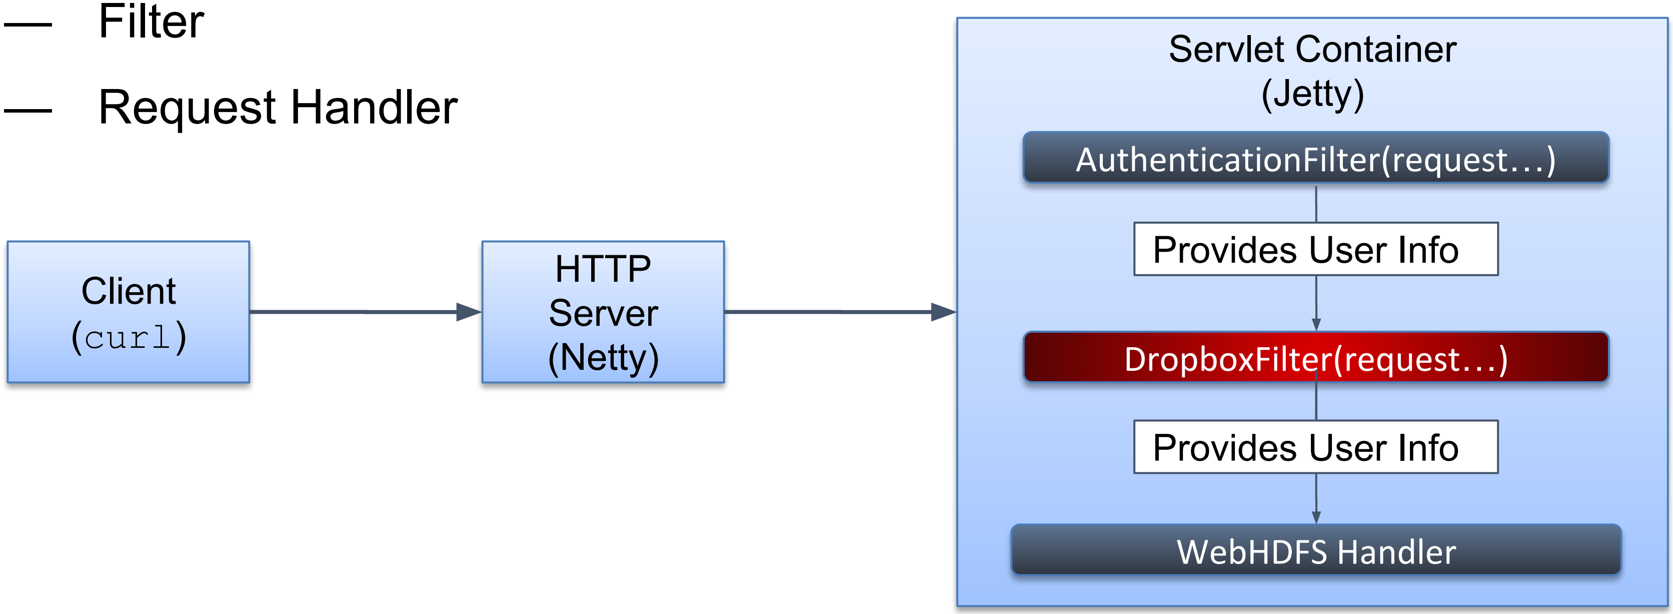

TechAtBloomberg.com

Bloomberg

Engineering

DropboxFilter for HDFS (Examples)

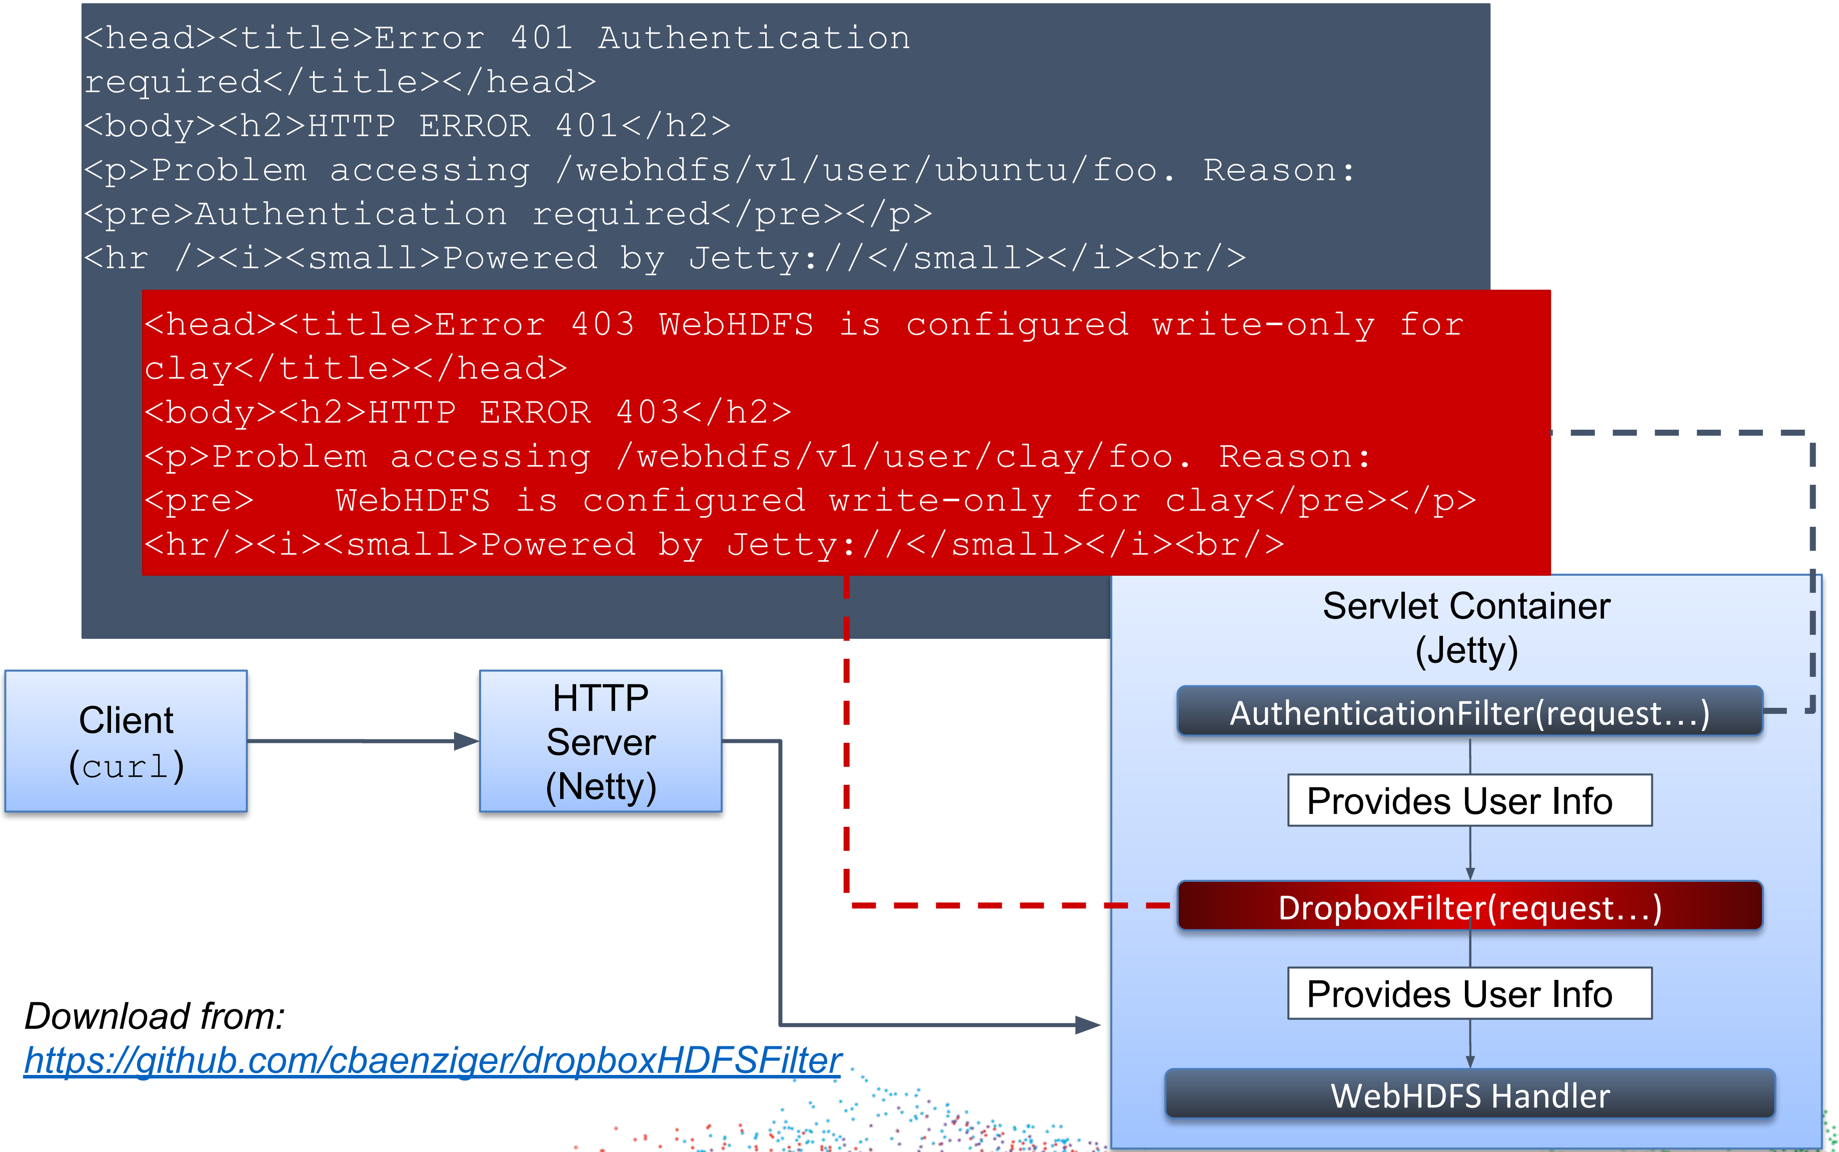

```
<head><title>Error 401 Authentication
required</title></head>
<body><h2>HTTP ERROR 401</h2>
<p>Problem accessing /webhdfs/v1/user/ubuntu/foo. Reason:
<pre>Authentication required</pre></p>
<hr /><i><small>Powered by Jetty://</small></i><br/>
```

```
<head><title>Error 403 WebHDFS is configured write-only for
clay</title></head>
<body><h2>HTTP ERROR 403</h2>
<p>Problem accessing /webhdfs/v1/user/clay/foo. Reason:
<pre>    WebHDFS is configured write-only for clay</pre></p>
<hr/><i><small>Powered by Jetty://</small></i><br/>
```

Download from:
https://github.com/cbaenziger/dropboxHDFSFilter

Plutonium Enclave

Image: Office of Legacy Management, U.S. D.O.E.,[CO-83-M-14 - Downdraft Table](), 20 Aug. 2014

Data Enclave

Centralized Models:

- Curator Model: Restrict access, operations and results
  - “The curator must remain present throughout the lifetime of the database”

    (Dwork, Cynthia. “Differential Privacy: A Survey of Results”, 1, Apr. 2008)
  - Statistical Disclosure Control
- Data Enclave:

  (Lane and Shipp. “Using a Remote Access Data Enclave for Data Dissemination”. *Intl. Jour. of Digital Curation*. 1.2 (2007))
  - Allow for Direct and Exact Access
  - Allow Arbitrary Computation

Isolation Glovebox

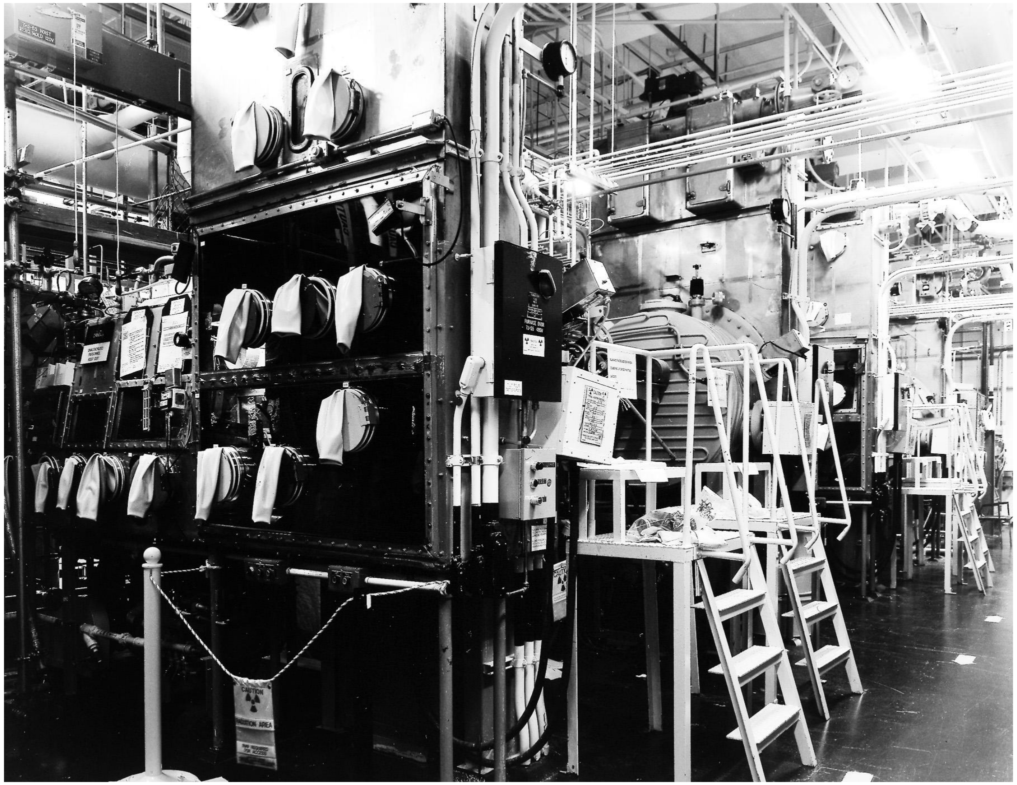

Images:
(Left) Library of Congress, U.S. D.O.E.,View Of A Worker Holding A Plutonium 'Button.' 19 Sep. 1973
(Right) Office of Legacy Management, U.S. D.O.E.,CO-83-M-8 - View of foundry induction furnaces. N.D.

Glovebox Production Line

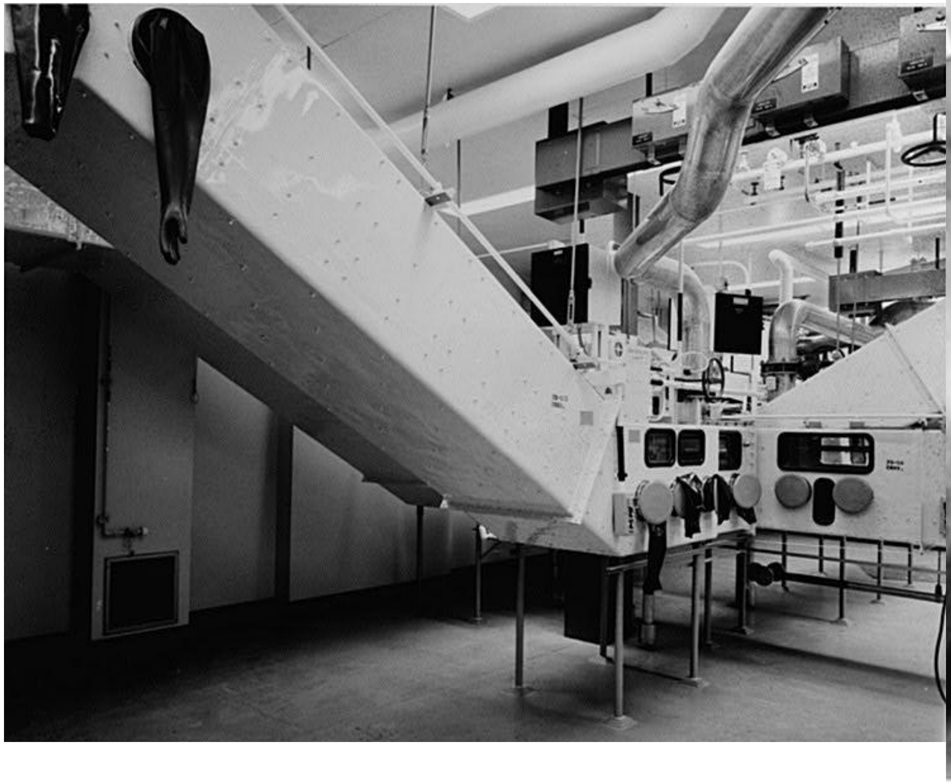

Images:
(Left) Library of Congress, U.S D.O.E., View Of A Glovebox Line Used In Plutonium Operations. 5 May. 1970
(Right) Office of Legacy Management, U.S. D.O.E., CO-83-M-3 - View of Chainveyor. 25 Jan. 1993

Data Glovebox

- **Leaded Pane of Glass**: Remote Desktop Without Download
- **Glove Ports:** Arbitrary Code Execution (Run code to manipulate)
- **Robotics:** Workflow Management
  - Deployment
  - Routine operations
- **Pass-throughs:**
  - Firewalls are insufficient
  - Protocol aware deep packet inspection
  - Databases
- **Firewalls:** Ensure user and workload isolation
  - Distributed file-systems
  - Local file-systems
  - Processing

Data Glovebox (Leaded Pane of Glass)

Avoid Overexposure to Raw Data

- Remote Desktop:
  - Limited RDP
- Key Attributes:
  - No copy out
  - No file shares
  - Isolation per user
- Useful to Have Tools:
  - Web browser for Jupyter/Zepplin
  - SSH client for command-line access

Data Glovebox (Glove Ports & Robotics)

Manipulate Your Data - With Code

- Run on a compute cloud using Apache YARN; submit:
  - SQL to Apache Hive
  - Python or Scala to Apache Spark
  - An arbitrary application
- Automation to ensure consistency (e.g. Apache Oozie)
  - A workflow manager for Hive and Spark jobs
  - Data transformations for expected reports -- known processes generating “decontaminated” results
  - Can run as a non-human service accounts to drop data in directory for data exfiltration
  - Can provide repeatable deployment of code using Git

Data Glovebox (Pass-Through)

Negative pressure (one-way) network; exfiltrate only “decontaminated” data

- Provide a process for data hand-off through an environment
- Firewalls:
  - Mostly a transport OSI Layer 4 device (TCP/IP)
  - Can do “deep packet inspection” - but need to MITM traffic
  - Policy rules for which users can manipulate which files become extensive
  - Prohibitively expensive
- DropboxFilter for HDFS
- Database RPCs are more complex but:
  - `GRANT INSERT ON DATABASE.* TO write_only@'%';`
  - `GRANT SELECT ON DATABASE.* TO read_only@'%';`

TechAtBloomberg.com

Bloomberg

Engineering

Firewalls (Workload and User Isolation)

Don’t let your data spontaneously combust; clean up “chips”

File Systems Leak

- Permission on data sets
- User collaboration locations
- Temporary/failed job data
- Temporary data locations

— Distributed file systems

— Hive Warehouse

— /tmp

— Local file systems

— /tmp, /var/tmp, /dev/shm

Take out the Trash

Unloading barrels of waste from a fire at Rocky Flats into Pit 10. The Department of Energy now considers Rocky Flats a triumph of the cleanup program. But much of the waste generated at the Rocky Flats site isn't cleaned up… it's here in Idaho. (69-6138)

Image: State of Idaho Oversight Monitor. Nov. 2006. Pg. 10

Private Temporary Directories

- To provide isolation one can use `pam_namespaces`
- Setup directories and clean-up one can use `pam_exec`

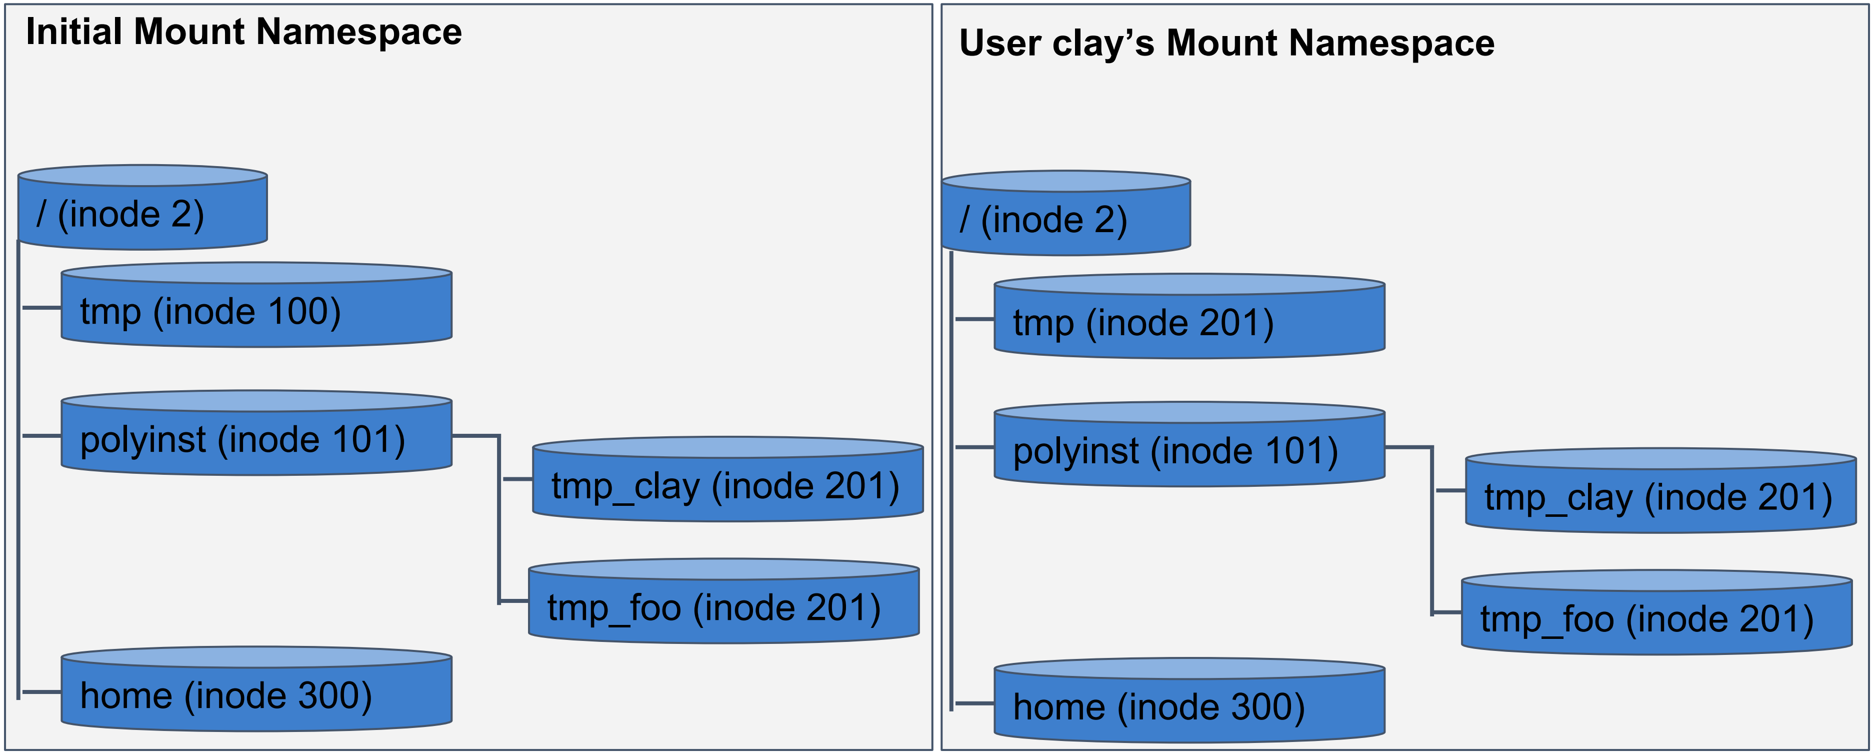

See also: Our integration of the work in https://github.com/bloomberg/chef-bach/pull/1278

Keeping the Pipes Flowing

Image: Office of Legacy Management, U.S. D.O.E., [CO-83-AF-1 - View of Building 215A.]() N.D.

YARN Network Isolation (Example)

YARN-7468 - Provide means for container network policy control

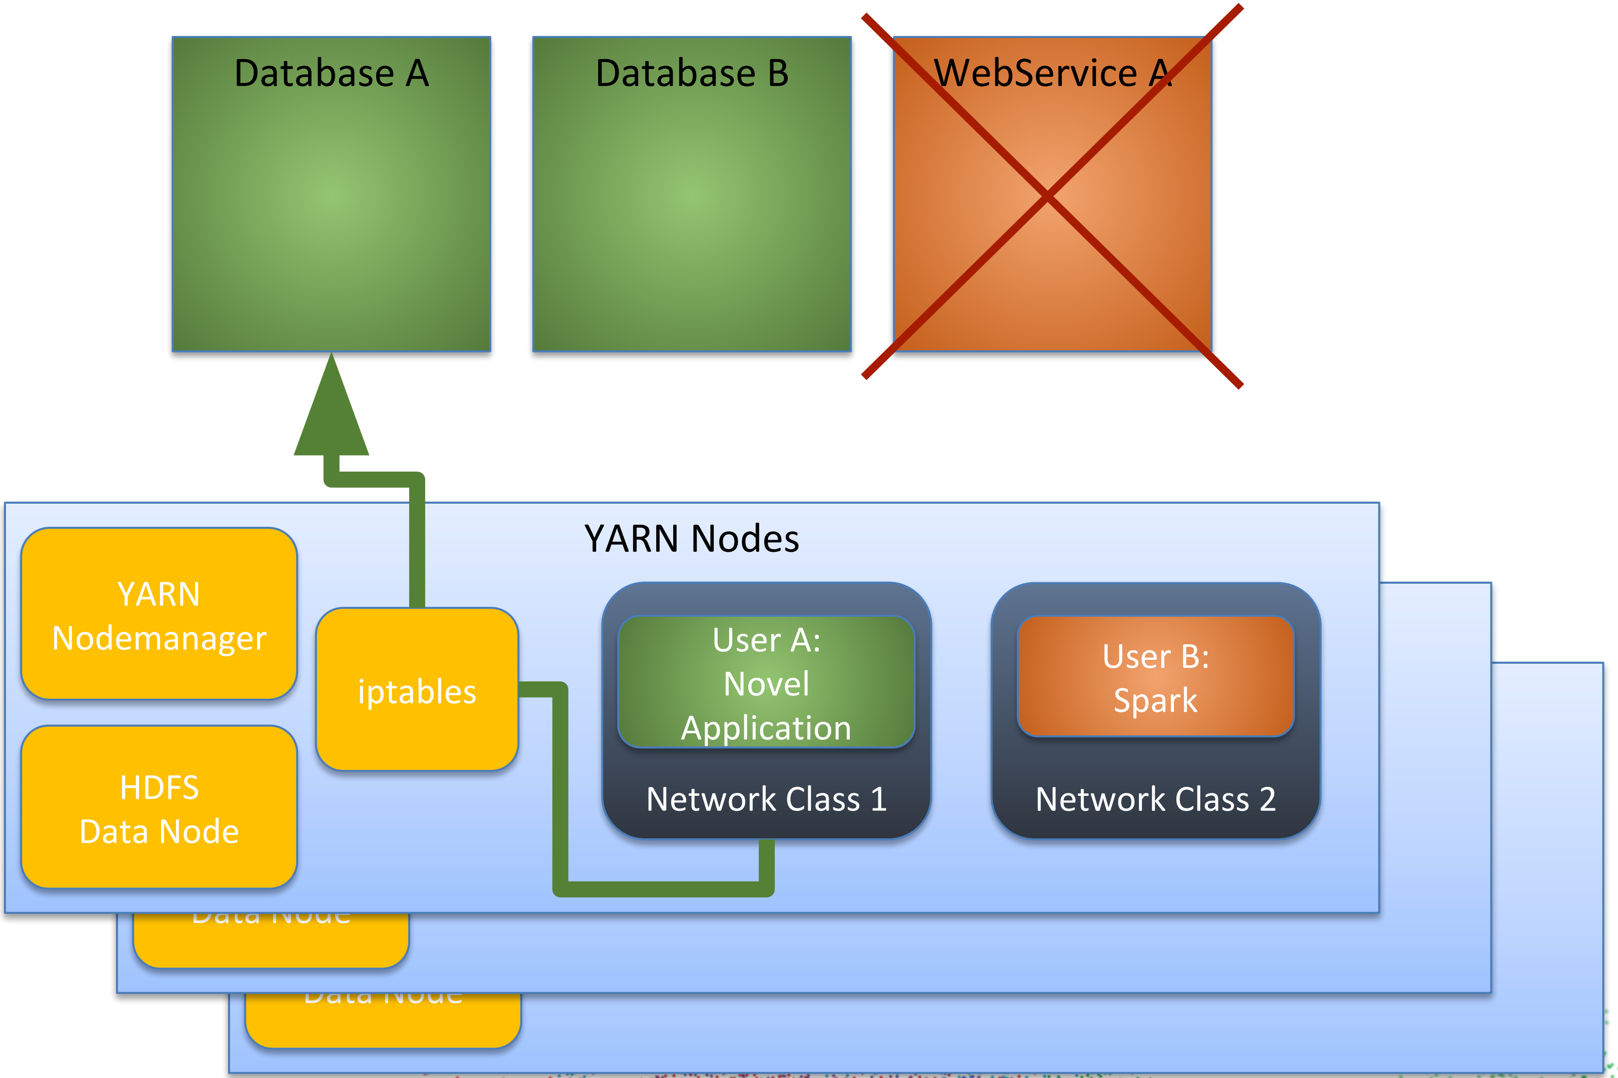

YARN Network Isolation (Example)

YARN-7468 - Provide means for container network policy control

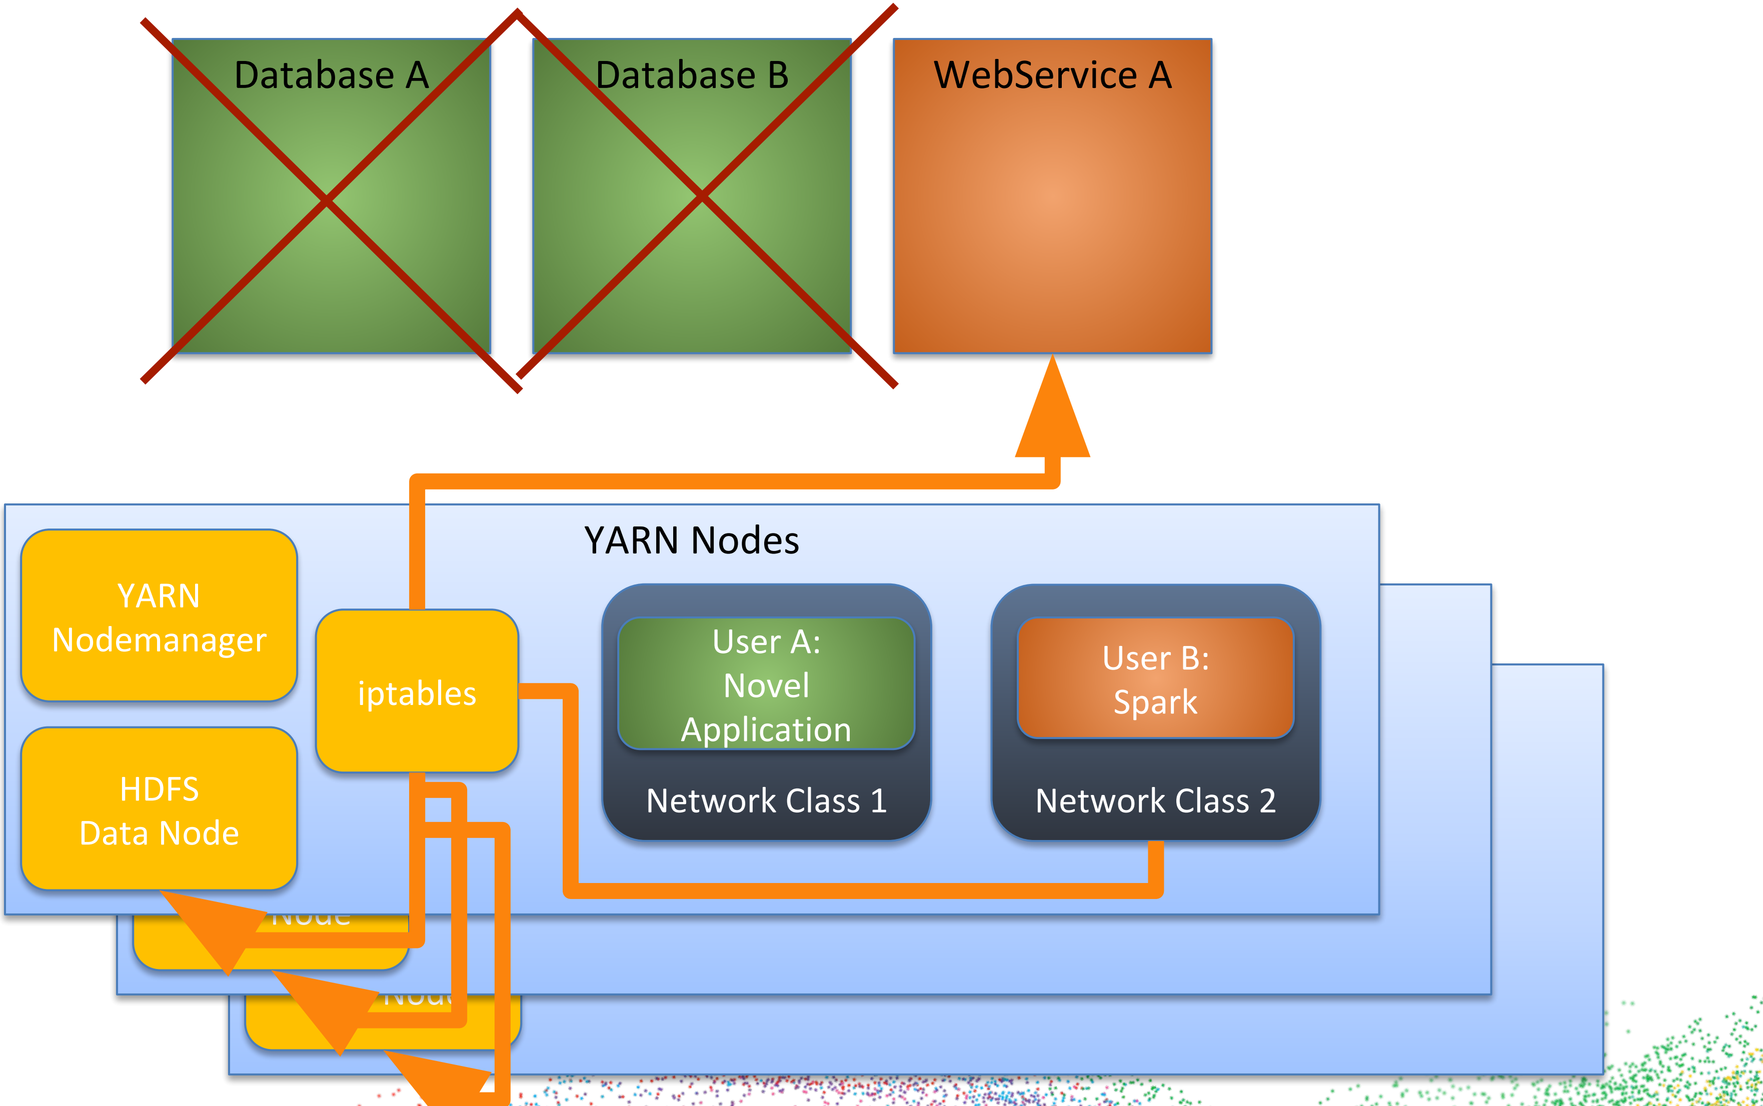

Firewalls Are Important

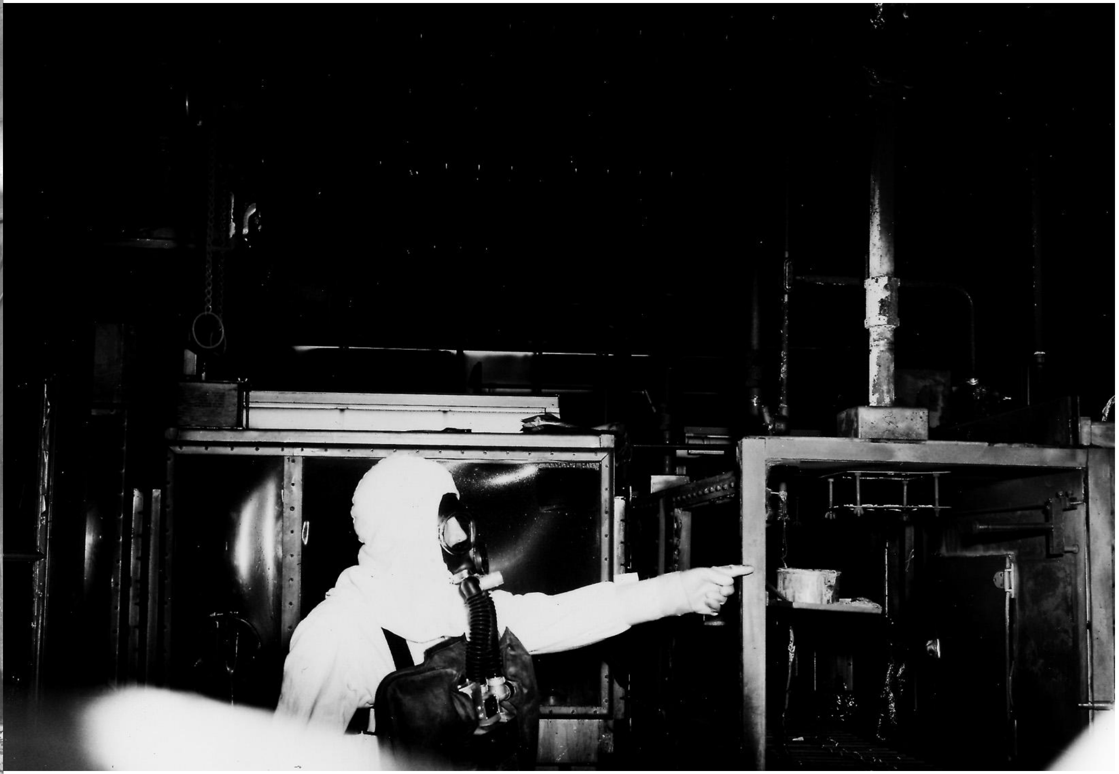

Image:
(Left) Office of Legacy Management, U.S. D.O.E., CO-83-N-3 - Damaged Filter Plenums. 16 Sept. 1957
(Right) Office of Legacy Management, U.S. D.O.E., CO-83-N-2 - Glove Box Where, on September 11, 1957, A Fire Started. 16 Sept. 1957

Data Glovebox

- **Leaded Pane of Glass**: Remote Desktop Without Copy
- **Glove Ports:** Manipulate your Data at An Arm's Length
- **Robotics:** Workflow Management
- **Pass-throughs:** Negative Pressure to Keep the Bits Flowing
- **Firewalls:** Ensure User and Workload Isolation

Thank You

Reference: http://github.com/bloomberg/chef-bach
Connect with Hadoop Team: hadoop@bloomberg.net

Bloomberg
Engineering

TechAtBloomberg.com